

1 **CLAIMS**

2 What is claimed is:

3 1. A method for building a data overlay, comprising:

4 providing a distributed hash table (DHT) that governs the insertion and retrieval  
5 of objects into and from a peer-to-peer system, wherein the distributed hash table  
6 includes a logical space including a plurality of DHT nodes having an associated plurality  
7 of DHT zones; and

8 building the data overlay as a data structure on top of the logical space of the  
9 distributed hash table by associating objects in the data structure with the DHT nodes,  
10 and by establishing links between the objects in the data structure.

11  
12 2. The method according to claim 1, wherein each link includes:

13 a first field that provides a hardwired pointer that points from a first object to a  
14 second object; and

15 a second field that provides a soft-state pointer that points from the first object to  
16 a DHT node which hosts the second object.

17  
18 3. The method according to claim 1, wherein the building of the data overlay  
19 makes use of:

20 a first primitive for setting a reference that establishes a pointer to an object in the  
21 distributed hash table;

22 a second primitive for returning an object referenced by a pointer; and

23 a third primitive for deleting an object referenced by a pointer.  
24  
25

1           4. The method according to claim 1, wherein the data overlay has a topology of a  
2 tree, the tree having a plurality of tree nodes associated with respective DHT nodes,  
3 wherein each tree node has a respective tree node zone associated therewith which  
4 corresponds to a part of the logical space of the distributed hash table.

5  
6           5. The method according to claim 4, wherein each tree node in the data overlay  
7 includes a key member which identifies a key associated with its respective tree node  
8 zone.

9  
10          6. The method according to claim 5, wherein the key has a value that is a function  
11 of coordinates that identify the center of the respective tree node zone.

12  
13          7. The method according to claim 4, wherein each tree node in the data overlay  
14 includes an operation member which defines an operation that is to be performed on data  
15 that is passed through the tree node.

16  
17          8. The method according to claim 4, wherein each tree node in the data overlay  
18 includes a report member which defines a report type that is to be generated using the tree  
19 node.

20  
21          9. The method according to claim 4, wherein the building of the data overlay  
22 comprises:

23           establishing a root tree node, the root tree node having a tree node zone  
24 corresponding to an entire span of the logical space of the distributed hash table.  
25

1           10. The method according to claim 4, wherein the building of the data overlay  
2 comprises:

3           examining a tree node zone associated with a particular tree node to determine  
4 whether the tree node zone is smaller than or equal to a DHT zone associated with the  
5 particular tree node's hosting DHT node; and

6           adding a child node associated with the particular tree node if the examining  
7 determines that the tree node zone is not smaller than or equal to the associated DHT  
8 zone.

9  
10          11. The method according to claim 10, further comprising repeating the  
11 examining and the adding for each tree node in the tree.

12  
13          12. A computer readable store including machine readable instructions for  
14 implementing the building of objects in the data overlay according to the method of claim  
15 10.

16  
17          13. A computer readable store having stored thereon a data overlay produced  
18 according to the method of claim 1.

19  
20          14. A computer readable store having stored thereon a data overlay having the  
21 topology of a tree produced according to the method of claim 4.

22  
23          15. A computer readable store having stored thereon a data structure, comprising:  
24  
25

1 a logical space of a distributed hash table (DHT), including a plurality of DHT  
2 nodes having a plurality of associated DHT zones, wherein the distributed hash table  
3 governs the insertion and retrieval of objects into and from a peer-to-peer system;

4 a data overlay implemented as a data structure on top of the logical space of the  
5 distributed hash table logical space, wherein the data overlay uses services provided by  
6 the distributed hash table in routing from one object to another in the data structure.

7  
8 16. The distributed computer readable store of claim 15, wherein the data overlay  
9 has a topology of a tree, the tree having a plurality of tree nodes associated with  
10 respective DHT nodes, wherein each tree node has a respective tree node zone associated  
11 therewith which corresponds to a part of the logical space of the distributed hash table.

12  
13 17. A method for passing data through a data overlay, comprising:  
14 providing a distributed hash table (DHT) that governs the insertion and retrieval  
15 of objects into and from a peer-to-peer system, wherein the distributed hash table  
16 includes a logical space including a plurality of DHT nodes having a plurality of  
17 associated DHT zones;

18 building a data overlay as a data structure on top of the logical space of the  
19 distributed hash table by associating objects in the data structure with the DHT nodes,  
20 and by establishing links between the objects in the data structure, wherein the data  
21 overlay defines a plurality of interconnected nodes; and

22 routing data through the data overlay by passing the data through its  
23 interconnected nodes.

1 . 18. The method according to claim 17, wherein the data overlay has a topology of  
2 a tree, the tree having a plurality of tree nodes associated with respective DHT nodes,  
3 wherein each tree node has a respective tree node zone associated therewith which  
4 corresponds to a part of the logical space of the distributed hash table.

5  
6 19. The method according to claim 18, wherein the routing of data through the  
7 data overlay includes gathering data from DHT nodes and passing the data up through the  
8 tree nodes to a root node of the tree.

9  
10 20. The method according to claim 18, wherein the routing of data through the  
11 data overlay includes disseminating data from a root node of the tree, through the tree  
12 nodes, to the DHT nodes.

13  
14 21. The method according to claim 18, wherein each tree node includes an  
15 operation member which defines an operation that is to be performed on data that is  
16 passed through the tree node.

17  
18 22. A computer readable store including machine readable instructions for  
19 implementing the routing of data through the data overlay according to the method of  
20 claim 17.

21  
22 23. A peer-to-peer system including a plurality of machines interacting in peer-to-  
23 peer fashion, comprising:

1 a logical space of a distributed hash table (DHT), including a plurality of DHT  
2 nodes having a plurality of associated DHT zones, wherein the distributed hash table  
3 governs the insertion and retrieval of objects into and from the peer-to-peer system; and

4 a data overlay implemented as a data structure on top of the logical space of the  
5 distributed hash table, wherein the data overlay uses services provided by the distributed  
6 hash table in routing from one object to another in the data structure,

7 wherein the logical space of the distributed hash table and the data overlay are  
8 implemented in distributed fashion in respective stores of the plurality of machines in the  
9 peer-to-peer system.

10  
11 24. The system according to claim 23, wherein the data overlay has a topology of  
12 a tree, the tree having a plurality of tree nodes associated with respective DHT nodes,  
13 wherein each tree node has a respective tree node zone associated therewith which  
14 corresponds to a part of the logical space of the distributed hash table.

15  
16 25. The system according to claim 24, further including routing logic configured  
17 to route data through the data overlay by passing the data through the tree nodes.

18  
19 26. The system according to claim 25, wherein the routing logic is configured to  
20 route the data through the data overlay by gathering data from DHT nodes and passing  
21 the data up through the tree nodes to a root node of the tree.

22  
23 27. The system according to claim 25, wherein the routing logic is configured to  
24 route data through the data overlay by disseminating data from a root node of the tree,  
25 through the tree nodes, to the DHT nodes.